

Objection to Drawings

The Office Action asserts that corrected drawings are required because Figures 1-7 are handwritten. Applicants submit herewith a set of formal drawings and respectfully request that any objection to the drawings be withdrawn.

Rejections Under 35 U.S.C. §102

Claims 1-24 are rejected under §102 as purportedly being anticipated by Pandya. These rejections are respectfully traversed.

I. Overview Of Embodiments Of The Invention

Embodiments of the invention are directed to a content addressable storage (CAS) system, which is one by which a unit of data stored on the CAS system is accessed using an address derived at least partially from the content of the unit of data. (specification, page 1, lines 10-12). The CAS system can serve as storage for a host computer. When a host computer sends a request to the CAS system to retrieve a unit of data, the host provides the content address of the unit of data, and the storage system then determines, based on the content address the physical location where the unit of data is stored so that it can be retrieved and returned to the host computer.

The task of determining the physical location for a unit of data may have several aspects, particularly when the storage system is a distributed storage system made up of a number of separate nodes. (page 1, lines 22-23). To determine the physical location of a unit of data on such a storage system, the storage system first determines on which node the unit of data is stored, and then determines which disk on that node the data is stored. (page 1, lines 26-28).

One known method of determining which storage nodes stores a particular unit of data is referred to as a multi cast location query (MLQ) (page 2, lines 9-10). Using this technique, a message is broadcast to each storage node that stores units of data, asking if it stores the particular unit of data to be accessed. (page 2, lines 10-13). Each storage node then determines if it stores the requested unit of data, and may do so by accessing a database or table that lists the units of data stored thereby. (page 2, lines 13-17). An MLQ is a computationally expensive

process, as it requires each storage node to perform an exhaustive database search for each unit of data requested. (page 3, lines 15-16).

To reduce the computational expense of using an MLQ to locate units of data on the storage system, another technique has been developed that employs a blob location index (BLI) to locate units of data (page 3, lines 20-23). The BLI is a database that maps the content addresses of units of data (“blobs”) to the storage nodes on which they are stored (page 3, lines 23-25). The administration of the BLI is split among the storage nodes, so that each storage node administers a portion thereof (page 3, lines 27-28). Thus, access requests for a unit of data need not be broadcast to all storage nodes, but just the one that administers the portion of the BLI that includes the requested unit of data. (page 3, lines 28-30). By distributing the BLI administration responsibility among the storage nodes, the computational burden of locating a particular unit of data is shared among the storage nodes. (page 4, line 31 – page 5, line 2).

In accordance with one embodiment of the invention, a content-addressable storage system is provided with an index (e.g., a BLI index) that maps a content address of at least one unit of data to a storage location at which it is stored, and the storage system maintains a cache of the location index to achieve performance improvements in accessing the information contained within the index. (page 12, lines 1-5).

The foregoing overview is provided merely to assist the Examiner in appreciating various aspects of the present invention. The overview may not apply to each of the independent claims, and the language of the independent claims may differ in material respects from the overview provided above. The Examiner is requested to give a careful consideration to the language of each of the independent claims and to address each on its own merits, without relying on the overview provided above. Applicants do not rely on the overview to distinguish any of the claims of the present invention over the prior art, but rather, rely only upon the arguments provided below.

II. Discussion of Pandya

Pandya is directed to a high performance hardware IP processor that is used to relieve the performance impact on a host processor of implementing the TCP/IP stack in software (Pandya,

¶0067). Figure 17 of Pandya is an architectural diagram of the IP processor (Pandya, ¶0107). Pandya discloses, in connection with Figure 17, that input queue block 1701 queues packets as they arrive and operates in conjunction with packet scheduler 1702, which retrieves packets from the input queue and passes them for classification to classification engine 1703 (Pandya, ¶0108). Classification engine 1703 receives a packet, classifies the packet, and provides a classification tag to the packet before it is provided to processor array 1706(a)...1706(n) (Pandya, ¶0108).

Figure 20 is more detailed diagram of classification engine 1703. Pandya discloses that the classification engine examines various fields of the received packet to identify the type of packet, the protocol type (e.g., IP, ICMP, TCP, UDP, etc.), the port addresses, the source and destination fields, and other fields (Pandya, ¶0118). Pandya discloses that the classifier, in some embodiments, may use a content addressable memory (CAM) array to assist in identifying packets (Pandya, ¶0118). As Applicants explained during the telephone interview, the classifier passes packet fields to the CAM to determine if the packet fields match any of the values stored in the CAM array (Pandya, ¶0118). When there is a match, the CAM provides an action/event tag. All of the action/event tags for a packet are compiled and are used to generate a classification tag for the packet, which may indicate the flow or session ID, the protocol type, or other classification information (Pandya, ¶0118).

III. Claims 1-8

Claim 1 is directed to a method of processing data in a computer system comprising at least one host and at least one content addressable storage system which stores data for the at least one host, wherein the at least one host accesses data units stored on the at least one storage system using content addresses generated based on the content of the data units. The method comprises: (a) maintaining at least one index that maps a content address of at least one data unit to a storage location within the at least one storage system at which the data unit is stored; and (b) maintaining a cache of the location index.

Pandya does not disclose or suggest a computer system comprising a content addressable storage system that stores data for at least one host computer. The Office Action states that Pandya discloses a content addressable storage system in the form of the content addressable

memory, disclosed in paragraph 0118, that forms a part of classification engine 1703 (*see* Office Action, page 2, ¶2). Applicants disagree that the content addressable memory is a “content addressable storage system which stores data for the at least one host,” as required by claim 1.

As discussed during the telephone interview, the content addressable memory of Pandya does not store data for a host computer. The content addressable memory of Pandya is a memory that is preprogrammed with packet fields and values (Pandya, ¶0118, lines 35-40). When network packets are received, data in the packets are compared to the values that are pre-stored in the content addressable memory so that the packet may be classified. **Pandya does not disclose or suggest that the content addressable memory stores packets provided by a host computer or retrieves packets from the content addressable memory and provides them to the host computer.**

Further, Pandya does not disclose or suggest that data units stored on the storage system (allegedly the CAM) are accessed using a content address generated based on the content of the data units. The Office Action asserts that Pandya discloses this feature at ¶0121, lines 76-100. However, the cited portion of Pandya does not even mention a content address, let alone a host accessing a data unit using the content address.

To be consistent with the rest of the rejection, the claimed content addressable storage system must purportedly read on the CAM array. However, as Applicants explained during the telephone interview, no matter what entity in Pandya is considered to be the host, Pandya does not disclose or suggest any entity accessing a data unit in the CAM array using a content address. The CAM array of Pandya serves as a lookup table which is used to classify a packet received from the host. The “address” used to access the CAM array is derived from the bits of the packet (Pandya, ¶0118). To the extent that it can be asserted that any “host computer” is accessing classification information stored in the CAM array, this information is not accessed using a content address computed from the content of the classification information stored in the content addressable memory. Pandya simply does not disclose or suggest a host computer accessing a data unit on a content addressable storage using a content address **that is computed from the content of the data unit.**

Additionally, Pandya does not disclose or suggest, “maintaining at least one index that maps a content address of at least one data unit to a storage location within the at least one storage system at which the data unit is stored” and “maintaining a cache of the location index.” The Office Action asserts that Pandya discloses maintaining an index that maps a content address of at least one data unit to a storage location at ¶0128, lines 11-14, which states, “[i]f there is no tag match in the cache array with the hash index, the look-up block uses this key to find the session entry from the external memory and replaces the current session cache entry with that session entry.” Thus, it appears that the Office Action asserts that the session cache disclosed by Pandya is an index that maps a content address of at least one data unit to a storage location at which the unit of data is stored. Applicants respectfully disagree with any such assertion. The session cache of Pandya shown at block 1704 of Figure 17 and in greater detail in Figure 29 is used to cache and store the TCP/IP session database and also the storage session database for a certain number of active sessions. **Nowhere does Pandya disclose that this hash index maps a content address of a data unit to a storage location for that data unit.**

The Office Action asserts that Pandya discloses maintaining a cache of the location index at ¶0128, lines 7-11, which states, “[t]he session cache look-up engine, block 2904, provides the functionality to look-up a specific session cache entry. This look-up block creates a hash index out of the fields provided or is able to accept a hash key and looks-up the session cache entry.” Applicants would like to clarify that the session cache in Pandya is not a cache for the hash index and the hash index is not a cache for the session cache. Rather, the session cache stores TCP/IP session information. The hash index is simply an index of the session cache that allows for look-up of a session cache entry.

Although it is unclear whether the Office Action interprets the hash index to be a cache of the session cache, whether the Office Action interprets the session cache to be a cache of the hash index, or whether the Office Action interprets the cited portion in some other way, there is simply no disclosure or suggestion that a cache of an index that maps a *content address* of a data unit to a storage location for that data unit is maintained.

In view of the foregoing, claim 1 patentably distinguishes over Pandya. Accordingly, it is respectfully requested that the rejection of claim 1 under 35 U.S.C. §102(e) be withdrawn.

Claims 2-8 depend from claim 1 and are patentable for at least the same reasons. Accordingly, it is respectfully requested that the rejection of these claims be withdrawn.

IV. Claims 9-16

Claim 9 is directed to at least one computer readable medium encoded with instructions that, when executed on a computer system, perform substantially the same method as recited in claim 1.

Thus, claim 9 patentably distinguishes over Pandya. Accordingly, it is respectfully requested that the rejection of claim 9 under 35 U.S.C. §102(e) be withdrawn.

Claims 10-16 depend from claim 9 and are patentable for at least the same reasons. Accordingly, it is respectfully requested that the rejection of these claims be withdrawn.

V. Claims 17-24

Claim 17 is directed to a content addressable storage system for use in a computer system, that includes the content addressable storage system and at least one host, wherein the at least one host accesses data units stored on the content addressable storage system using content addresses generated based on the content of the data unit. The content addressable storage system comprises: at least one storage device to store data received from the at least one host; and at least one controller that: maintains at least one index that maps a content address of at least one data unit to a storage location within the content addressable storage system at which the data unit is stored; and maintains a cache of the location index.

As should be clear from the discussion above, Pandya fails to disclose or suggest a content addressable storage system, wherein a host accesses data units stored on the content addressable storage system using content addresses generated based on the content of the data unit, and fails to disclose or suggest a controller that maintains at least one index that maps a content address of at least one data unit to a storage location within the content addressable storage system at which the data unit is stored and maintains a cache of the location index.

Thus, claim 17 patentably distinguishes over Pandya. Accordingly, it is respectfully requested that the rejection of claim 17 under 35 U.S.C. §102(e) be withdrawn.

Claims 18-24 depend from claim 17 and are patentable for at least the same reasons.
Accordingly, it is respectfully requested that the rejection of these claims be withdrawn.

CONCLUSION

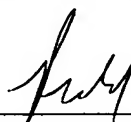

In view of the foregoing amendments and remarks, this application should now be in condition for allowance. A notice to this effect is respectfully requested. If the Examiner believes, after this amendment, that the application is not in condition for allowance, the Examiner is requested to call the Applicant's attorney at the telephone number listed below.

If this response is not considered timely filed and if a request for an extension of time is otherwise absent, Applicant hereby requests any necessary extension of time. If there is a fee occasioned by this response, including an extension fee, that is not covered by an enclosed check, please charge any deficiency to Deposit Account No. 23/2825.

Dated: April 11, 2006

Respectfully submitted,

Stephen J. Todd, et. al., Applicants

By  

Richard F. Giunta

Registration No.: 36,149

WOLF, GREENFIELD & SACKS, P.C.

Federal Reserve Plaza

600 Atlantic Avenue

Boston, Massachusetts 02210-2206

(617) 646-8000